



06-04-07

AF
JRW

PATENT

Attorney Docket No. 07447.0061-00

Xerox Ref. No.: D/AOA25

CERTIFICATE UNDER 37 CFR § 1.10 OF MAILING BY "EXPRESS MAIL"

EV860818916US

May 31, 2007

USPS Express Mail Label Number

Date of Deposit

I hereby certify that this correspondence is being deposited with the United States Postal Services "Express Mail Post Office to Addressee" service under 37 CFR §1.10 on the date indicated above and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

By: Cindy Bagletto
Cindy Bagletto

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

In re Application of:)
Eric GAUSSIER et al.) Group Art Unit: 2161
Application No.: 09/982,236) Examiner: NGUYEN, Cam Linh T.
Filed: October 19, 2001)
For: METHODS, SYSTEMS AND) Confirmation No.: 7611
ARTICLES OF MANUFACTURE)
FOR SOFT hIERARCHICAL)
CLUSTERING OF CO-)
OCCURRING OBJECTS)

Mail Stop Appeal Brief--Patents

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

TRANSMITTAL OF SUPPLEMENTAL APPEAL BRIEF (37 C.F.R. 41.37, 41.39)

Transmitted herewith is the SUPPLEMENTAL APPEAL BRIEF in this application with respect to the Office Action mailed March 1, 2007 and the Notice of Appeal originally filed on November 22, 2005.

This application is on behalf of

Small Entity Large Entity

Pursuant to 37 C.F.R. 41.20(b)(2), the fee for filing the Appeal Brief is:

- \$250.00 (Small Entity)
- \$500.00 (Large Entity)
- This fee was previously paid on May 19, 2006

TOTAL FEE DUE:

Notice of Appeal Fee \$0.00 (previously paid on November 22, 2005)

Extension Fee (if any) \$0.00

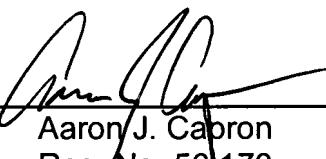
Total Fee Due \$0.00

- Enclosed is a check for \$[Text] to cover the above fees.

PETITION FOR EXTENSION. If any extension of time is necessary for the filing of this Appeal Brief, and such extension has not otherwise been requested, such an extension is hereby requested, and the Commissioner is authorized to charge necessary fees for such an extension to Deposit Account 06-0916.

FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER, L.L.P.

Dated: May 31, 2007

By: 
Aaron J. Capron
Reg. No. 56,170

Post Office Address (to
which correspondence is
to be sent)

Finnegan, Henderson, Farabow,
Garrett & Dunner, L.L.P.
901 N.Y. Ave. N.W.
Washington, D.C. 20001-4413
(202) 408-4000



PATENT
Customer Number 22,852
Attorney Docket No. 07447.0061-00000
Xerox Ref. No.: D/AOA25

CERTIFICATE UNDER 37 CFR § 1.10 OF MAILING BY "EXPRESS MAIL"
EV860818916US May 31, 2007

USPS Express Mail Label Number

Date of Deposit

I hereby certify that this correspondence is being deposited with the United States Postal Services "Express Mail Post Office to Addressee" service under 37 CFR §1.10 on the date indicated above and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

By: Cindy Baglietto

Cindy Baglietto

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

In re Application of:)
Eric GAUSSIER et al.)
Application No.: 09/982,236) Group Art Unit: 2161
Filed: October 19, 2001) Examiner: NGUYEN, Cam Linh T.
For: METHODS, SYSTEMS AND ARTICLES) Confirmation No.: 7611
OF MANUFACTURE FOR SOFT)
HIERARCHICAL CLUSTERING OF CO-)
OCCURRING OBJECTS)

Mail Stop Appeal Brief--Patents

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450.

Sir:

SUPPLEMENTAL APPEAL BRIEF UNDER 37 C.F.R. §§ 41.37, 41.39

In response to the Office Action mailed on March 1, 2007, and pursuant to 37 C.F.R. §§ 41.37, 41.39, Appellants present the following Supplemental Appeal Brief and request reinstatement of the Appeal, which was originally filed on November 22, 2005.

The Notice of Appeal and the fee set forth in 37 C.F.R. § 41.37 were originally filed on November 22, 2005 along with a request for a pre-appeal brief conference. The final rejection was withdrawn and prosecution was reopened following the decision of the pre-appeal brief conference panel. The Examiner then issued another final rejection

on Dec. 21, 2005, which was appealed in the Appeal Brief filed May 19, 2006. In response to the Appeal Brief, Examiner withdrew the previous final rejection and entered new grounds of rejection in an Office Action mailed March 1, 2007 ("Office Action"). Appellants appeal Examiner's rejection of the claims and request that the Board of Appeals reverse in whole the rejections of claims 1 - 26 and order the allowance of these claims.

I. Real Party Interest

The real party in interest is Xerox Corporation.

II. Related Appeals and Interferences

There are no related appeals or interferences at this time.

III. Status of Claims

Claims 1-26 are pending and stand rejected. Appellants reproduce claims 1-26 in Claims Appendix for the Board's convenience.

IV. Status of Amendments

All amendments for this application have been entered.

V. Summary of Invention

The application describes methods, systems, and articles of manufacture for soft hierarchical clustering of objects based on a co-occurrence of object pairs. Clustering allows data to be hierarchically grouped (or clustered) based on its characteristics, so that objects, such as text data in documents, similar to each other are placed in a common cluster in a hierarchy. In soft hierarchical clustering, an object may be assigned to more than one cluster in a hierarchy, as opposed to a hard assignment, whereby an object is assigned to only one cluster in the hierarchy.

A modified Expectation-Maximization (EM) process is performed on object pairs reflecting documents and words, respectively, such that a given class of the objects ranges over all nodes of a topical hierarchy (as opposed to the leaves alone) and the assignment of a document to a topic may be based on any ancestor of the given class. Moreover, the assignment of a given document to any topic in the hierarchy may also be based on a particular (document, word) pair under consideration during the process. The modified EM process may be performed for every child class generated from an ancestor class until selected constraints associated with the topical hierarchy are met. A representation of the resultant hierarchy of topical clusters may be created and made available to entities that request the topics of the document collection. See, e.g., pg. 4, lines 22-23; and pg. 5, lines 1-11.

The modified algorithm eliminates the reliance on leaf nodes alone and allows any set S_i to be explained by a combination of any leaves and/or ancestor nodes included in an induced hierarchy. That is, i objects may not be considered as blocks, but rather as pieces that may be assigned in a hierarchy based on any j co-occurring objects. In one configuration, a topical clustering application performed by a computer may assign parts of a document i to different nodes in an induced hierarchy for different words j included in the document i . See, e.g., pg. 15, lines 10-20.

For example, the probability of observing any pair of co-occurring objects, such as documents and words (i, j) , may be modeled by defining a variable I_{α} (controls the assignment of documents to a hierarchy) such that it is dependent on the particular document and word pair (i, j) under consideration during a topical clustering process. In one configuration, the class α may range over all nodes in an induced hierarchy to

assign a document (i object) to any node in the hierarchy, not just leaves. Furthermore, by defining a class v as any ancestor of α in the hierarchy, the nodes may be hierarchically organized. See, e.g., pg. 15, lines 21-23; and pg. 16, lines 1-6.

Different j objects may be generated from different vertical paths of an induced hierarchy; that is, from paths in the hierarchy associated with non null values of $I_{i\alpha}$. Furthermore, because α may be any node in the hierarchy, the i objects may be assigned to different levels of the hierarchy. Accordingly, implementation of the model results in a pure soft hierarchical clustering of both i and j objects by eliminating any hard assignments of these objects. See, e.g., pg. 18, lines 10- 21.

The model may be implemented for a variety of applications, depending upon the meaning given to objects i and j . For example, it may be applied to document clustering based on topic detection. In such a configuration, i objects may represent documents and j objects may represent words included in the documents. Clusters or topics of documents may be represented by leaves and/or nodes of an induced hierarchy. The topics associated with the document collection may be obtained by interpreting any cluster as a topic defined by the word probability distributions, $p(j|v)$. The soft hierarchical model may take into account several properties when interpreting the clusters, such as: (1) a document may cover (or be explained by) several topics (soft assignment of i objects provided by the probability $p(i|\alpha)$); (2) a topic best described by a set of words, which may belong to different topics due to polysemy (the property of a word to exhibit several different, but related meanings) and specialization (soft assignment of j objects provided by the probability $p(j|v)$); and (3) topics may be

hierarchically organized, which corresponds to the hierarchy induced over clusters.

See, e.g., pg. 20, lines 25-30; and pg. 21, lines 1-11.

One or more conditions associated with a hierarchy that may be induced may allow a computer to determine when an induced hierarchy reaches a desired structure with respect to the clusters defined therein. For example, a condition may be defined that instructs a processor to stop locating co-occurring objects (i, j) in a document collection that is being clustered based on a predetermined number of leaves, and/or a level of the induced hierarchy. See, e.g., pg. 23, lines 1-11.

Pending independent claim 1 recites a method performed by a computer for clustering a plurality of documents in a structure comprised of a plurality of clusters hierarchically organized, wherein each document includes a plurality of words and is represented as a set of (document, word) pairs. See, e.g., pg. 2, lines 11-18; pg. 19, lines 12-18; Fig. 5; pg. 20, lines 13-23; and Fig. 6. The method comprises: accessing the document collection; performing a clustering process that creates a hierarchy of clusters that reflects a segregation of the documents in the collection based on the words included in the documents, wherein any document in the collection may be assigned to a first cluster in the hierarchy based on a first segment of the respective document, and the respective document may be assigned to a second cluster in the hierarchy based on a second segment of the respective document, wherein the first and second clusters are associated with different paths of the hierarchy. See, e.g., pg. 2, lines 19-23; pg. 3, lines 1-4; pg. 20, lines 1-8; Fig. 5; pg. 21, lines 18-23; pg. 22, lines 1-22; pg. 23, lines 1-23; pg. 24, lines 1-23; pg. 25, lines 1-7; Fig. 6; pg. 29, lines 17-18; pg. 30 lines 1-2; and Fig. 7. A representation of the hierarchy of clusters is stored in

memory and made available to an entity in response to a request associated with the document collection. See, e.g., pg. 20, lines 13-23; and pg. 21, lines 1-17. Claims 2 - 7 all ultimately depend from claim 1.

Pending independent claim 8 recites a method performed by a computer for determining topics of a document collection, by accessing the document collection, wherein each document includes a plurality of words and is represented as a set of (document, word) pairs. See, e.g., pg. 2, lines 11-18; pg. 19, lines 12-18; Fig. 5; pg. 20, lines 1-8 and 13-19; and Fig. 6. The method comprises performing a clustering process including: creating a tree of nodes that represent topics associated with the document collection based on the words in the document collection, wherein any node in the tree may include a word that is shared by another node in the tree, and assigning fragments of one or more documents included in the document collection to multiple nodes in the tree based on the (document, word) pairs and storing a representation of the tree in a memory. See, e.g., pg. 20, lines 1-8 and 13-19; pg. 21, lines 18-23; pg. 22, lines 1-22; pg. 23, lines 1-23; pg. 24, lines 1-23; pg. 25, lines 1-23; pg. 26, lines 1-5; Fig. 6; pg. 29, lines 17-18; pg. 30 lines 1-2; and Fig. 7. The representation is made available for processing operations associated with the document collection. See, e.g., pg. 20, lines 13-23; and pg. 21, lines 1-17. Claim 9 ultimately depends from claim 8.

Pending independent claim 10 recites a method performed by a processor for clustering data in a database by receiving a collection of documents, wherein each document includes a plurality of words and is represented as a set of (document, word) pairs. See, e.g., pg. 2, lines 11-18; pg. 19, lines 12-18; Fig. 5; pg. 20, lines 13-23; pg. 25, lines 18-23; pg. 26, line 1; and Fig. 6. The method comprises creating a first

ancestor node reflecting a first topic based on words included in the collection of documents; creating descendant nodes from the first ancestor node, each descendant node reflecting descendant topics based on the first node, until a set of leaf nodes reflecting leaf topics are created. See, e.g., pg. 14, lines 15-22; pg. 15, lines 1-13; pg. 16, lines 13-20; pg. 19, lines 12-18; and Fig. 5. The step of creating descendant nodes includes assigning each document in the collection to a plurality of descendant and leaf nodes; and providing a set of topics associated with the collection of documents based on the created nodes and assignment of documents, wherein the descendant and leaf nodes may be created based on one or more words included in more than one document in the collection of documents. See, e.g., pg. 19, lines 12-18; pg. 20, lines 1-8; pg. 21, lines 18-23; pg. 22, lines 1-4 and 11-22; pg. 23, lines 1-23; pg. 24, lines 1-23; pg. 25, lines 1-23; pg. 26, lines 1-5; Fig. 6; pg. 26, lines 6-22; pg. 27, lines 1-16; pg. 28, lines 19-22; pg. 29, lines 1-3 and 17-18; pg. 30 lines 1-2; and Fig. 7. Claim 11 ultimately depends from claim 10.

Pending independent claim 12 recites a method performed by a processor for clustering data in a database by receiving a collection of documents, wherein each document includes a plurality of words and is represented as a set of (document, word) pairs. See, e.g., pg. 2, lines 11-18; pg. 4, lines 16-23; pg. 5, lines 1-4; pg. 7, lines 22-23; Fig. 1; pg. 19, lines 12-18; Fig. 5; pg. 20, lines 13-23; pg. 25, lines 18-23; pg. 26, line 1; and Fig. 6. The method comprises creating a hierarchy of nodes based on the words in the collection of documents, each node reflecting a topic associated with the documents, wherein the hierarchy of nodes includes ancestor nodes, descendant nodes, and leaf nodes. Each document in the collection is assigned to a plurality of

nodes in the hierarchy, wherein each document may be assigned to any of the ancestor, descendant, and leaf nodes. See, e.g., pg. 14, lines 15-22; pg. 15, lines 3-13; pg. 16, lines 13-20; pg. 19, lines 12-18; and Fig. 5. A set of topic clusters associated with the collection of documents is provided and based on the created nodes and assignment of documents, wherein the hierarchy may include a plurality of nodes that are each created based on a same set of words included in the collection of documents. See, e.g., pg. 26, lines 6-22; pg. 27, lines 1-16; pg. 28, lines 19-22; pg. 29, lines 1-3 and 17-18; pg. 30 lines 1-2; and Fig. 7.

Pending independent claim 13 recites a method performed by a computer for clustering data stored on a computer-readable medium by receiving a collection of data objects, represented as a set of (first data object, second data object) pairs. See, e.g., pg. 2, lines 11-18; pg. 4, lines 16-23; pg. 5, lines 1-4 and 18-20; pg. 6, lines 19-23; pg. 7, lines 1-3, 12-16 and 22-23; Fig. 1; pg. 19, lines 12-18; Fig. 5; pg. 20, lines 13-23; pg. 25, lines 18-23; pg. 26, line 1; and Fig. 6. The method comprises: for each first data object: assigning the first data object to a first node in a hierarchy of nodes based on the second data objects included in the first data object, wherein the first node may be any node included in the hierarchy and wherein two or more nodes in the hierarchy may share the same second object; creating a final hierarchy of nodes arranged in clusters based on the assignment of the first data objects. See, e.g., pg. 20, lines 1-8 and 13-19; pg. 21, lines 18-23; pg. 22, lines 1-22; pg. 23, lines 1-23; pg. 24, lines 1-23; pg. 25, lines 1-23; pg. 26, lines 1-5; Fig. 6; pg. 29, lines 17-18; pg. 30, lines 1-2 and 11-15; and Fig. 7. A representation of the final hierarchy is stored in memory and made available

to an entity in response to a request associated with the collection of first data objects.

See, e.g., pg. 20, lines 13-23; and pg. 21, lines 1-17.

Pending independent claim 14 recites a method performed by a processor for clustering data in a database by receiving a request from a requesting entity to determine topics associated with a collection of documents, each document including a plurality of words and being represented as a set of (document, word) pairs. See, e.g., pg. 2, lines 11-18; pg. 19, lines 12-18; Fig. 5; pg. 20, lines 13-23; pg. 25, lines 18-23; pg. 26, line 1; and Fig. 6. The method comprises determining the topics associated with the collection of documents based on a hierarchy including a plurality of clusters, wherein each cluster reflects a topic and a document in the collection may be assigned to a set of clusters in the hierarchy based on different words included in the document, and wherein each cluster in the set may be associated with different paths in the hierarchy. See, e.g., pg. 2, lines 19-23; pg. 3, lines 1-4; pg. 20, lines 1-8; Fig. 5; pg. 21, lines 18-23; pg. 22, lines 1-22; pg. 23, lines 1-23; pg. 24, lines 1-23; pg. 25, lines 1-7; Fig. 6; pg. 29, lines 17-18; pg. 30 lines 1-2; and Fig. 7. A representation of the hierarchy is stored in memory and made available to the requesting entity. See, e.g., pg. 20, lines 13-23; and pg. 21, lines 1-17.

Pending independent claim 15 recites a computer-implemented method for clustering a plurality of multi-word documents into a hierarchical data structure including a root node associated with a plurality of sub-nodes, wherein each sub-node is associated with a topic cluster based on the plurality of documents. See, e.g., pg. 2, lines 11-18; pg. 7, lines 22-23; pg. 19, lines 12-18; Fig. 5; pg. 20, lines 1-8 and 13-19; and Fig. 6. The method comprises: retrieving a first document; associating the first

document with a first topic cluster based on a first portion of the first document; associating the first document with a second topic cluster based on a second portion of the document; and providing a representation of topics associated with the plurality of multi-word documents based on the hierarchical data structure including the first and second topic clusters, wherein the first and second topic clusters are associated with a different sub-node. See, e.g., pg. 14, lines 15-22; pg. 15, lines 3-13; pg. 16, lines 13-20; pg. 19, lines 12-18; Fig. 5; pg. 25, lines 7-23; pg. 26, lines 1-5; Fig. 6; pg. 26, lines 6-22; pg. 27, lines 1-16; pg. 28, lines 19-22; pg. 29, lines 1-3 and 17-18; pg. 30, lines 1-2; and Fig. 7. Claims 16 - 19 all ultimately depend from claim 15.

Pending independent claim 20 recites a computer-implemented method for clustering data reflecting users, represented as a set of (data, user) pairs, into a hierarchical data structure including a root node associated with a plurality of sub-nodes, wherein each sub-node represents an action that is performed on a document collection. See, e.g., pg. 31, lines 1-4 and 22-23; pg. 30, lines 11-15; pg. 2, lines 11-18; pg. 7, lines 22-23; pg. 19, lines 12-18; Fig. 5; pg. 20, lines 1-8 and 13-19; and Fig. 6.

The method comprises: accessing a user data collection reflecting a plurality of users who each perform at least one action on the document collection, wherein each action may be unique; performing a clustering process that creates the hierarchical data structure, wherein the clustering processing comprises: retrieving a first user data, associated with a first user, from the user data collection, associating the first user data with a first sub-node based on a first action performed by the first user on the document collection, and associating the first user data with a second sub-node provided the first user data is based on a second action, wherein the first and second sub-nodes are

associated with different descendent paths of the hierarchical data structure. See, e.g., pg. 31, lines 1-4 and 22-23; pg. 30, lines 11-15; pg. 2, lines 19-23; pg. 3, lines 1-4; pg. 20, lines 1-8; Fig. 5; pg. 21, lines 18-23; pg. 22, lines 1-22; pg. 23, lines 1-23; pg. 24, lines 1-23; pg. 25, lines 1-7; Fig. 6; pg. 29, lines 17-18; pg. 30, lines 1-2; and Fig. 7. A representation of the hierarchical data structure is stored in memory and made available to an entity in response to a request associated with the user data collection. See, e.g., pg. 20, lines 13-23; and pg. 21, lines 1-17. Claim 21 ultimately depends from claim 20.

Pending independent claim 22 recites a computer-implemented method for clustering a plurality of images based on text associated with the images, where each image is represented as a set of pairs (image, image feature) and (image, text feature), into a hierarchical data structure including a root node associated with a plurality of sub-nodes, wherein each sub-node represents a different topic. See, e.g., pg. 31, lines 5-7 and 22-23; pg. 30, lines 11-15; pg. 2, lines 11-18; pg. 7, lines 22-23; pg. 19, lines 12-18; Fig. 5; pg. 20, lines 1-8 and 13-19; and Fig. 6. The method comprises: accessing an image collection; performing a clustering process that creates the hierarchical data structure, wherein the clustering processing comprises: associating a first image with a first sub-node based on a first portion of text associated with the first image, and associating the first image with a second sub-node based on a second portion of text associated with the first image, wherein the first and second sub-nodes are associated with different descendant paths of the hierarchical data structure. See, e.g., pg. 31, lines 5-7 and 22-23; pg. 30, lines 11-15; pg. 2, lines 19-23; pg. 3, lines 1-4; pg. 20, lines 1-8; Fig. 5; pg. 21, lines 18-23; pg. 22, lines 1-22; pg. 23, lines 1-23; pg. 24, lines 1-23; pg. 25, lines 1-7; Fig. 6; pg. 29, lines 17-18; pg. 30, lines 1-2; and Fig. 7. A

representation of the hierarchical data structure is stored in memory and made available to an entity in response to a request associated with the image collection. See, e.g., pg. 20, lines 13-23; and pg. 21, lines 1-17.

Pending independent claim 23 recites a computer-implemented method for clustering customer purchases, represented as a set of (customer, purchase) pairs, into a hierarchical data structure including a root node associated with a plurality of sub-nodes, wherein each sub-node represents a group of customers who purchased the same type of product from one or more business entities. See, e.g., pg. 31, lines 8-23; pg. 30, lines 11-15; pg. 2, lines 11-18; pg. 7, lines 22-23; pg. 19, lines 12-18; Fig. 5; pg. 20, lines 1-8 and 13-19; and Fig. 6. The method comprises: accessing information associated with a plurality of customers who purchased various types of products from a plurality of business entities; performing a clustering process that creates the hierarchical data structure, wherein the clustering processing comprises: associating a first customer with a first sub-node based on a first type of product purchased from a first business entity, and associating the first customer with a second sub-node provided the first customer is based on a second type of product that the first customer purchased from a second business entity, wherein the first and second sub-nodes are associated with different descendant paths of the hierarchical data structure. See, e.g., pg. 31, lines 8-23; pg. 30, lines 11-15; pg. 2, lines 19-23; pg. 3, lines 1-4; pg. 20, lines 1-8; Fig. 5; pg. 21, lines 18-23, pg. 22, lines 1-22; pg. 23, lines 1-23; pg. 24, lines 1-23; pg. 25, lines 1-7; Fig. 6; pg. 29, lines 17-18; pg. 30, lines 1-2; and Fig. 7. A representation of the hierarchical data structure is stored in memory and made available in response to a request associated with the customer data collection. See, e.g., pg.

20, lines 13-23; and pg. 21, lines 1-17. Claims 24-26 all ultimately depend from claim 23.

VI. Grounds of Rejection to Be Reviewed on Appeal

A. Whether claims 1- 26 should be rejected under 35 U.S.C. § 103 as unpatentable in light of U.S. Patent No. 6,742,003 ("Heckerman") in view of U.S. Patent No. 6,460,025 ("Fohn").

VII. Argument

The case law and the MPEP set forth the requirements to establish a *prima facie* case of obviousness, and both place the burden of doing so squarely on the examiner. Specifically, the Examiner must meet three basic criteria. First, the prior art reference (or references when combined) must teach or suggest all claim limitations. *MPEP* § 2142. Second, there must be some reason why a person of ordinary skill in the art would have combined the prior art elements in the manner claimed. *USPTO Memorandum* from Margaret A. Focarino, Deputy Commissioner for Patent Operations, May 3, 2007, page 2. Third, there must be a reasonable expectation of success. *MPEP* § 2142. Appellants respectfully assert that the Examiner has not met the burden of establishing one or more of these basic requirements.

Within the Office Action, the Examiner rejected claims 1-26 by using Heckerman in view of Fohn. The Examiner alleged that Heckerman disclosed all of the limitations of claim 1 except for the limitation "wherein any document in the collection may be assigned to a first cluster in the hierarchy based on a first segment of the respective document, and the respective document may be assigned to a second cluster in the hierarchy based on a second segment of the respective document." Office Action at

page 3. Contrary to the specification of Heckerman, the Examiner attempted to minimize Heckerman's deficiency and even insinuated that Heckerman discloses this missing limitation by stating that "Heckerman does not clearly disclose [the missing limitation.] Heckerman only mentions that the document has n attributes (col. 27, line 67), and based on the matches or those attribute settings, a document can belong to multiple clusters in the hierarchical tree and therefore, forming a multi-level hierarchical organization." Office Action at pages 3-4. By minimizing Heckerman's deficiency, it appears that the Examiner is justifying combining the weak teaching of Fohn with Heckerman. But, as provided in the analysis below, Heckerman is quite clear that a document can only belong to one cluster.

Heckerman is directed to a visualization scheme for the graphical depiction of clusters and cluster relationships. Specifically, Heckerman discloses methods for reviewing web page access patterns of web users to optimize links between various web pages or to customize advertisements to the demographics of the users. Heckerman permits the visualization of clusters and cluster relationships between the web pages based on user access patterns where the data is represented as a collection of records, each containing values for various attributes. See Heckerman at col. 1, lines 27 - 46.

Figures 1A - 1D provide an overview of the methods outlined in Heckerman. Fig. 1B illustrates the results of the manual **classification** of a collection of records shown in Fig. 1A. **Classification** techniques allow the manual grouping of the records of a collection into classes. Once classification has been performed, a new record may be automatically classified when it is added to the collection, as shown in Fig. 1C. In

contrast, **clustering** techniques provide an automated process for analyzing the records of the collection and identifying clusters of records that have similar attributes. Figure 1D illustrates the results of the clustering process performed in Heckerman on the collection of records shown in Fig.1A. Heckerman at col. 1, line 47 - col. 2, line 23. As stated in Heckerman, "the values stored in the column marked "CLUSTER" in FIG. 1D have been determined by the clustering algorithm." *Id.* at col. 2, lines 15 -18 and lines 27 - 32. The clustering method described in Heckerman only assigns a record to one of several clusters. *Id.* at col. 2, lines 28 - 32.

Other sections in Heckerman also reiterate that a record may only belong to one cluster. For example, Heckerman states:

Clustering process 1510 automatically, and using a conventional clustering process, such as "EM" clustering, reads, as symbolized by lines 1503, data for the cases, in a dataset (population or collection) stored within case data 100 and automatically determines applicable **mutually exclusive** categories for these cases and then categorizes (classifies) each of those cases into those categories.

Id. at col. 25, lines 9 - 15 (emphasis added).

Further, the requirement that clusters contain mutually exclusive records underpins techniques outlined in Heckerman. For example, the technique to compute a discriminative score for cluster (group) c1 versus cluster (group) c2 given observation X=x, requires that clusters c1 and c2 are mutually exclusive. *Id.* at col. 32, lines 57 - 59.

Other passages in Heckerman also teach that a record may only belong to one cluster.

For each such user, database 1360 contains dataset 100 that contains a record for each such user along with predefined attributes (illustratively numbered 1 through j) for that user, and **the class (category or cluster)** to which that record is categorized. As noted, each such record together with all its attributes is commonly referred to as a "case". In addition,

database 1360 also contains cluster data 1355 which specifies, e.g., clusters, segment and segment hierarchies.

Id. at col. 21, lines 43 - 51 (emphasis added).

Moreover, as clearly indicated in Fig. 18 and the associated description, the sum of the percentages of cases in individual segments is equal to the total population. If a case belonged to multiple clusters (and therefore to multiple segments), then the sum of the percentage of cases in individual segments would exceed the total population. According to Heckerman, “[s]egments are clusters of cases that exhibit similar behavior, such as users on a given site, and have similar properties, such as age or gender. A segment consists of a summary of the database records (cases) that belong to it.” *Id.* at col. 21, lines 61 - 64. Heckerman further states:

Display 1800 shows segment hierarchy 1810 in a left portion of the display. A user, such as a business manager or data analyst, by clicking on a down arrow displayed within hierarchy 1810 can expand a segment group to expose its constituent segments, as shown. Each segment and group are listed along with their corresponding percentages of an entire population. In that regard, segment 5 represents 10% of the entire population, segment group 6 represents 27% of the entire population, and so forth. As depicted, segment group 6 also contains segments 3 and 4.

Id. at col. 22, lines 28 - 38.

As shown in Fig. 18, the individual percentages of the total number of cases associated with segments 5, 6, and 8, are 10%, 27%, and 63%, respectively, which add up to 100%, representing the total population.

Indeed, the clustering process in Heckerman automatically reads data for the cases in a dataset, automatically determines applicable mutually exclusive categories (e.g., classes and clusters) for these cases, and then categorizes each of those cases into those categories. Because categories contain mutually exclusive cases, a case cannot belong to more than one cluster. Claims 1, 19, 38, and 53 in Heckerman also

clearly recite that the data records are classified based on the attribute / value pairs associated with each such record, into a plurality of “**mutually exclusive** first clusters” (emphasis added) further refuting the Examiner’s opinion. *Id.* at col. 33, lines 50-53. Thus, contrary to the Examiner’s downplaying of Heckerman’s deficiencies, Heckerman is quite clear that a data record can only belong to one cluster.

In an attempt to overcome Heckerman’s deficiency, the Examiner provided Fohn as a secondary reference. Fohn provides a forward-looking indication of relevant user-selectable nodes (which represent categories and sub-categories) during the navigation and browsing of pre-categorized collections. Fohn at 8:19-24, 60-65. In particular, Fohn provides a computer software program including exploration capabilities for a category hierarchy by using categories. *Id.* at 7:9-14. The categories used to organize entities (products, parts, or other entities of, e.g., a catalog) correspond to a hierarchical association that, in general, can be described in three steps. *Id.* at 7:15-20. First, structural relevance of each entity is calculated for each node in a particular hierarchy based on the existence of associations between a node and any entity. *Id.* at 8:25-31. Second, state relevance of an entity is calculated for the current state of the solution, where the solution state is based upon the previously visited nodes and the current node. *Id.* at 8:31-35. Third and finally, entity relevance (structural and state relevance) is combined with node-based forward-checking to provide a hierarchical exploration scheme to guide user’s exploration over multiple node hierarchies and their entities. *Id.* at 8:35-38. This solution requires two necessary pieces of information: (1) a pre-existing hierarchy of category nodes and (2) a pre-existing base of entities (collection of entities) already instances of the category nodes. *Id.* at 8:39-56.

A. Heckerman in view of Fohn fail to disclose each and every limitation of the claims

MPEP § 2143.03 requires that “[t]o establish *prima facie* obviousness of a claimed invention, all of the claim limitations must be taught or suggested by the prior art.” MPEP § 2143.03 (citing *In re Royka*, 490 F.2d 981 (C.C.P.A. 1974)) (emphasis added). According to this section, “[a]ll words in a claim must be considered in judging the patentability of that claim against the prior art.” *Id.* (citing *In re Wilson*, 424 F.2d 1382, 1385 (C.C.P.A. 1970)) (emphasis added). For claims 1-26 to be obvious, Heckerman in view of Fohn must disclose each and every claim limitation.

Claim 1 is directed to a method performed by a computer, the method including “performing a clustering process that creates a hierarchy of clusters that reflects a segregation of the documents in the collection based on the words included in the documents, wherein any document in the collection may be assigned to a first cluster in the hierarchy based on a first segment of the respective document, and the respective document may be assigned to a second cluster in the hierarchy based on a second segment of the respective document, wherein the first and second clusters are associated with different paths of the hierarchy ...” (hereinafter known as “performing step”).

The Examiner rejected the performing step in an interesting manner. First, the Examiner rejected the limitations “performing a clustering process that creates a hierarchy of clusters that reflects a segregation of the documents in the collection based on the words included in the documents” and “wherein the first and second clusters are associated with different paths of the hierarchy “ by using Heckerman. Office Action at page 3. The Examiner then extracted out of the performing step the limitation “wherein

any document in the collection may be assigned to a first cluster in the hierarchy based on a first segment of the respective document, and the respective document may be assigned to a second cluster in the hierarchy based on a second segment of the respective document.” The Examiner acknowledged that Heckerman did not disclose this extracted portion and then alleged that Fohn overcame Heckerman’s deficiencies by disclosing this extracted portion. *Id.* at 3-4. MPEP § 2141.02 recites “[i]n determining the differences between the prior art and the claims, the question under 35 U.S.C. [§] 103 is not whether the differences themselves would have been obvious, but whether the claimed invention as a whole would have been obvious.” According to the MPEP, the Examiner must look at the performing step as a whole when determining whether Heckerman in view of Fohn discloses the entire performing step. But the Examiner failed to consider the performing step as a whole when applying this combination of references.

The Examiner attempted to support the rejection by stating that “[c]learly Fohn teaches that an entity can be placed in two different categories.” Office Action at 4. But the Examiner over-simplifies the differences because the claimed invention requires that the documents used to create the hierarchy of clusters are the same documents that can be assigned to both first and second clusters of the hierarchy of clusters. As acknowledged by the Examiner, Heckerman fails to disclose, teach, or suggest the ability to assign documents to both first and second clusters based on segments within the same documents used to create the hierarchy of clusters. Office Action at pages 3-4. *In arguendo*, assuming that the Examiner’s allegations regarding Fohn are true, at most, Fohn teaches an entity being associated with multiple nodes and does not

disclose creating a hierarchy of clusters by assigning these entities to multiple nodes.

So like Heckerman, Fohn fails to teach, suggest, or disclose the ability to assign documents to both first and second clusters based on segments within the same documents used to create the hierarchy of clusters.

Further, the Examiner failed to tie in how Fohn's entities are segregated based on words included in the entities. Just because an entity points to both a Group Portrait and a Birthday node, as shown in Fig. 6A of Fohn, that does not necessarily mean that the entity had those words within the documents themselves. For example, a person—not a computer—simply could have segregated these entities based either on a product specification or on their own personal knowledge, but not necessarily on the segments in the document itself.

Furthermore, the performing step requires a computer performing a clustering process that creates a hierarchy of clusters—based on words in a collection of documents—having a first cluster and second cluster, within different paths of the hierarchy, that provides a document associated with each cluster. But the Examiner assumes that just because an entity can be assigned to different nodes, without disclosing how the hierarchy is created, whether it be by a person, etc., that this shows how a computer performs a clustering process that creates a hierarchy of clusters—having first and second clusters, each associated with a document—that reflects the segregation of the documents in the collection based on the words in the documents. In fact, Fohn requires a pre-existing relationship between the nodes and entities before determining whether the nodes and the entities are relevant. Because Fohn requires a pre-existing relationship, Fohn cannot supplement Heckerman's creation of a hierarchy

of clusters based on a collection of documents when a document of this collection can be applied to both a first and second cluster of the hierarchy being created.

For at least these reasons, Heckerman in view of Fohn fails to disclose, teach, or suggest claim 1 as a whole.

Claim 8 recites “[a] method performed by a computer for determining topics of a document collection, the method comprising: ...performing a clustering process including: creating a tree of nodes that represent topics associated with the document collection based on the words in the document collection, wherein any node in the tree may include a word that is shared by another node in the tree, and assigning fragments of one or more documents included in the document collection to multiple nodes in the tree based on the (document, word) pairs ...” (emphasis added).

It appears that the Examiner is relying on Fohn to disclose this limitation because the Examiner did not explicitly address any limitation of claim 8 in the Office Action. The Examiner stated, with regards to claim 1, that “[c]learly Fohn teaches that an entity can be placed in two different categories.” Office Action at 4. But Fohn does not assign entities to nodes because these nodes and entities already have a pre-existing relationship. *Id.* at 8:39-56. Accordingly, because of this pre-existing relationship, Fohn could not further assign entities to multiple nodes based on a pair involving the entity. The Examiner may attempt to rely on Heckerman to illustrate the assigning, but Heckerman only assigns a record to one of several clusters. For at least these reasons, Heckerman in view of Fohn fails to disclose each and every limitation of claim 8.

B. One of ordinary skill in the art would not have a reason to combine Fohn into Heckerman.

For establishing the motivation to combine Fohn into Heckerman, the Examiner stated that:

It would have obvious to one of ordinary skill in the art at the time the invention was made to apply the teaching of Fohn into the invention of Heckerman because the combination would 'provide a powerful flexible technique for locating entities in a large information space using hierarchical navigation and browsing of these or more hierarchies'. (Col. 24, lines 14 - 17 of Fohn). The combination system would [enable] a user to search for a solution meeting his selected constraints from a multi-perspective viewpoint, guiding him through ascent and descent in a hierarchy as well as lateral exploration and movement to other hierarchies (col. 24, lines 19-23 of Fohn).

Office Action at page 4.

As provided in the analysis above, Heckerman is quite clear that, when creating a hierarchy, a data record can only belong to one cluster. Because Heckerman provides a data record that can only belong to one cluster, Heckerman fails to disclose the claimed invention. To overcome Heckerman's deficiency, the Examiner alleged that Fohn, which discloses a single entity belonging to multiple nodes, could be combined into Heckerman. But by doing so completely contradicts the invention disclosed in Heckerman. Therefore, one of ordinary skill in the art would not have a reason to combine Fohn into Heckerman to form the claimed invention.

In arguendo, even if one of ordinary skill in the art would find a reason to combine Fohn into Heckerman, the combination would produce a result that would be completely different from the claims. Heckerman discloses creating a hierarchy having a cluster, wherein a data record can only belong to one cluster. As described above, Fohn discloses determining the structural and state relevance of a pre-existing collection of entities that are already instances of the pre-existing category nodes. Fohn

is quite clear that it does not create hierarchies. *Id.* at 21:4-8. This would limit the applicability of Fohn to an established hierarchy. If combining Fohn with Heckerman, based on the motivation provided by the Examiner, one of ordinary skill in the art would derive a combination involving Heckerman's creation of a hierarchy having a cluster, wherein a data record can only belong to one cluster and then applying Fohn's intelligent exploration to Heckerman's hierarchy to determine if that data record—only associated to the one cluster—was relevant to the one cluster that the data record was associated with. This combined invention is far different from the claims at issue.

Further, the Examiner has failed to establish why one of ordinary skill in the art would even combine pre-categorized classification of Fohn with Heckerman's clustering that involves no categorization hierarchy having data record only associated with one cluster.

For at least these reasons, one of ordinary skill in the art would not have a reason to combine Heckerman and Fohn for the limitations provided in claim 1. Accordingly, because Heckerman in view of Fohn fails to disclose each and every limitation of claim 1 and because one of ordinary skill in the art would not have a reason to combine Heckerman and Fohn for the matter claimed, the Examiner has failed to establish a *prima facie* obviousness regarding claim 1. Therefore, Appellants respectfully submit that claim 1 is patentable over the cited prior art.

Claims 2-7 and 24-26 depend on claim 1 and are patentable for at least the same reasons as is claim 1.

For the reasons outlined above, the Examiner has failed to establish a *prima facie* case of obviousness regarding the process of:

performing a clustering process including: creating a tree of nodes that represent topics associated with the document collection based on the words in the document collection, wherein any node in the tree may include a word that is shared by another node in the tree, and assigning fragments of one or more documents included in the document collection to multiple nodes in the tree based on the (document, word) pairs;

as recited in claim 8. Accordingly, Appellants respectfully submit that claim 8 is patentable over the cited prior art. Claim 9 depends on claim 8 and is patentable for at least the same reasons as is claim 8.

For the reasons outlined above, the Examiner has failed to establish a *prima facie* case of obviousness regarding the process of:

assigning each document in the collection to a plurality of descendant and leaf nodes; and providing a set of topics associated with the collection of documents based on the created nodes and assignment of documents, wherein the descendant and leaf nodes may be created based on one or more words included in more than one document in the collection of documents

as recited in claim 10. Accordingly, Appellants respectfully submit that claim 10 is patentable over the cited prior art. Claim 11 depends on claim 10 and is patentable for at least the same reasons as is claim 10.

For the reasons outlined above, the Examiner has failed to establish a *prima facie* case of obviousness regarding the process of:

assigning each document in the collection to a plurality of nodes in the hierarchy, wherein each document may be assigned to any of the ancestor, descendant, and leaf nodes; and providing a set of topic clusters associated with the collection of documents based on the created nodes and assignment of documents, wherein the hierarchy may include a plurality of nodes that are each created based on a same set of words included in the collection of documents

as recited in claim 12. Accordingly, Appellants respectfully submit that claim 12 is patentable over the cited prior art.

For the reasons outlined above, the Examiner has failed to establish a *prima facie* case of obviousness regarding the process of:

assigning the first data object to a first node in a hierarchy of nodes based on the second data objects included in the first data object, wherein the first node may be any node included in the hierarchy and wherein two or more nodes in the hierarchy may share the same second object; creating a final hierarchy of nodes arranged in clusters based on the assignment of the first data objects

as recited in claim 13. Accordingly, Appellants respectfully submit that claim 13 is patentable over the cited prior art.

For the reasons outlined above, the Examiner has failed to establish a *prima facie* case of obviousness regarding the process of:

determining the topics associated with the collection of documents based on a hierarchy including a plurality of clusters, wherein each cluster reflects a topic and a document in the collection may be assigned to a set of clusters in the hierarchy based on different words included in the document, and wherein each cluster in the set may be associated with different paths in the hierarchy

as recited in claim 14. Accordingly, Appellants respectfully submit that claim 14 is patentable over the cited prior art.

For the reasons outlined above, the Examiner has failed to establish a *prima facie* case of obviousness regarding the process of:

providing a representation of topics associated with the plurality of multi-word documents based on the hierarchical data structure including the first and second topic clusters, wherein the first and second topic clusters are associated with a different sub-node

as recited in claim 15. Accordingly, Appellants respectfully submit that claim 15 is patentable over the cited prior art. Claims 16-19 depend on claim 15 and are patentable for at least the same reasons as is claim 15.

For the reasons outlined above the Examiner has failed to establish a *prima facie* case of obviousness regarding the process of:

associating the first user data with a second sub-node provided the first user data is based on a second action, wherein the first and second sub-nodes are associated with different descendant paths of the hierarchical data structure

as recited in claim 20. Accordingly, Appellants respectfully submit that claim 20 is patentable over the cited prior art. Claim 21 depends on claim 20 and is patentable for at least the same reasons as is claim 20.

For the reasons outlined above the Examiner has failed to establish a *prima facie* case of obviousness regarding the process of:

associating the first image with a second sub-node based on a second portion of text associated with the first image, wherein the first and second sub-nodes are associated with different descendant paths of the hierarchical data structure

as recited in claim 22. Accordingly, Appellants respectfully submit that claim 22 is patentable over the cited prior art.

For the reasons outlined above, the Examiner has failed to establish a *prima facie* case of obviousness regarding the process of:

associating the first customer with a second sub-node provided the first customer is based on a second type of product that the first customer purchased from a second business entity, wherein the first and second sub-nodes are associated with different descendant paths of the hierarchical data structure

as recited in claim 23. Accordingly, Appellants respectfully submit that claim 23 is patentable over the cited prior art.

VIII. Conclusion

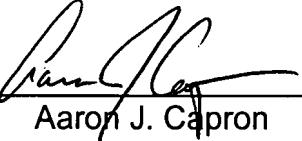
For the foregoing reasons, Appellants respectfully request reversal of all of the bases for rejection set forth in the Grounds of Rejection to Be Reviewed on Appeal section above (i.e., Section VI.A) and allowance of all pending claims.

Please grant any extensions of time required to enter this paper and charge any additional required fees to our Deposit Account No. 06-0916.

Respectfully submitted,

FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER, L.L.P.

Dated: May 31, 2007

By: 
Aaron J. Capron
Reg. No. 56,170

Post Office Address (to
which correspondence is
to be sent)

Finnegan, Henderson, Farabow,
Garrett & Dunner, L.L.P.
901 N.Y. Ave. N.W.
Washington, D.C. 20001-4413
(202) 408-4000

Claims Appendix

Pending claims on appeal:

1. A method performed by a computer for clustering a plurality of documents in a structure comprised of a plurality of clusters hierarchically organized, wherein each document includes a plurality of words and is represented as a set of (document, word) pairs, the method comprising:

accessing the document collection;

performing a clustering process that creates a hierarchy of clusters that reflects a segregation of the documents in the collection based on the words included in the documents, wherein any document in the collection may be assigned to a first cluster in the hierarchy based on a first segment of the respective document, and the respective document may be assigned to a second cluster in the hierarchy based on a second segment of the respective document, wherein the first and second clusters are associated with different paths of the hierarchy;

storing a representation of the hierarchy of clusters in a memory; and

making the representation available to an entity in response to a request associated with the document collection.

2. The method of claim 1, wherein performing a clustering process comprises:

assigning the document collection to a first class;

setting a probability parameter to an initial value; and

determining, for each document in the collection at the value of the parameter, a probability of an assignment of the document in the collection to a cluster in the hierarchy based on a word included in the document and the first class.

3. The method of claim 2, wherein the step of determining further comprises: determining whether the first class has split into two child classes, wherein each child class reflects a cluster descendant from an initial cluster reflected by the first class; and

increasing the value of the parameter based on the determination whether the first class has split into two child classes.

4. The method of claim 3, further comprising: repeating the step of determining, for each document in the collection at the value of the parameter, and the step of increasing the value of the parameter until the first class has split into two child classes.

5. The method of claim 4, further comprising: performing the clustering process for each child class until each of the respective child class splits into two new child classes reflecting clusters descendant from the respective child class.

6. The method of claim 5, further comprising: repeating the clustering process for each new child class such that a hierarchy of clusters is created, until a predetermined condition associated with the hierarchy is met.

7. The method of claim 6, wherein the predetermined condition is one of a maximum number of leaves associated with the hierarchy and depth level of the hierarchy.

8. A method performed by a computer for determining topics of a document collection, the method comprising:

accessing the document collection, each document including a plurality of words and being represented as a set of (document, word) pairs;

performing a clustering process including:

creating a tree of nodes that represent topics associated with the document collection based on the words in the document collection, wherein any node in the tree may include a word that is shared by another node in the tree, and

assigning fragments of one or more documents included in the document collection to multiple nodes in the tree based on the (document, word) pairs;

storing a representation of the tree in a memory; and

making the representation available for processing operations associated with the document collection.

9. The method of claim 8, wherein the step of assigning comprises:

associating a set of documents in the document collection with a first class reflecting all of the nodes in the tree, wherein the set of documents may include all or some of the documents in the collection;

defining a second class reflecting any ancestor node of a node in the first class;

determining, for each document in the set, a probability that different words included in a respective document co-occurs with the respective document in any node in the tree based on the first and second classes; and

assigning one or more fragments of any document in the set to any node in the tree based on the probability.

10. A method performed by a processor for clustering data in a database, the method comprising:

receiving a collection of documents, each document including a plurality of words and being represented as a set of (document, word) pairs;

creating a first ancestor node reflecting a first topic based on words included in the collection of documents;

creating descendant nodes from the first ancestor node, each descendant node reflecting descendant topics based on the first node, until a set of leaf nodes reflecting leaf topics are created,

wherein creating descendant nodes includes:

assigning each document in the collection to a plurality of descendant and leaf nodes; and

providing a set of topics associated with the collection of documents based on the created nodes and assignment of documents,

wherein the descendant and leaf nodes may be created based on one or more words included in more than one document in the collection of documents.

11. The method of claim 10, wherein the step of creating descendant nodes comprises:

selecting a first document in the collection;

defining a first class that includes all of the nodes;

defining a second class that may include any ancestor node of any node included in the first class; and

determining, for each document in the collection, a target word of an object pair including a target document and the target word such that the first document equals the target document in the object pair based on a probability associated with the first and second classes; and

assigning the first document to any ancestor, descendant, and leaf node based on the determining.

12. A method performed by a processor for clustering data in a database, the method comprising:

receiving a collection of documents, each document including a plurality of words and being represented as a set of (document, word) pairs;

creating a hierarchy of nodes based on the words in the collection of documents, each node reflecting a topic associated with the documents, wherein the hierarchy of nodes includes ancestor nodes, descendant nodes, and leaf nodes;

assigning each document in the collection to a plurality of nodes in the hierarchy, wherein each document may be assigned to any of the ancestor, descendant, and leaf nodes; and

providing a set of topic clusters associated with the collection of documents based on the created nodes and assignment of documents,

wherein the hierarchy may include a plurality of nodes that are each created based on a same set of words included in the collection of documents.

13. A method performed by a computer for clustering data stored on a computer-readable medium, the method comprising:

receiving a collection of data objects, represented as a set of (first data object, second data object) pairs;

for each first data object:

assigning the first data object to a first node in a hierarchy of nodes based on the second data objects included in the first data object, wherein the first node may be any node included in the hierarchy and wherein two or more nodes in the hierarchy may share the same second object;

creating a final hierarchy of nodes arranged in clusters based on the assignment of the first data objects;

storing a representation of the final hierarchy in a memory; and

making the representation of the final hierarchy available to an entity in response to a request associated with the collection of first data objects.

14. A method performed by a processor for clustering data in a database, the method comprising:

receiving a request from a requesting entity to determine topics associated with a collection of documents, each document including a plurality of words and being represented as a set of (document, word) pairs;

determining the topics associated with the collection of documents based on a hierarchy including a plurality of clusters, wherein each cluster reflects a topic and a document in the collection may be assigned to a set of clusters in the hierarchy based on different words included in the document, and wherein each cluster in the set may be associated with different paths in the hierarchy;

storing a representation of the hierarchy in a memory; and

making the representation available to the requesting entity.

15. A computer-implemented method for clustering a plurality of multi-word documents into a hierarchical data structure including a root node associated with a plurality of sub-nodes, wherein each sub-node is associated with a topic cluster based on the plurality of documents, the method comprising:

retrieving a first document;

associating the first document with a first topic cluster based on a first portion of the first document;

associating the first document with a second topic cluster based on a second portion of the document; and

providing a representation of topics associated with the plurality of multi-word documents based on the hierarchical data structure including the first and second topic clusters,

wherein the first and second topic clusters are associated with a different sub-node.

16. The method of claim 15, wherein the first and second portions contain at least one unique word.

17. The method of claim 15, wherein associating the first document with a first topic cluster comprises:

assigning the plurality of multi-word documents to a first class;

setting a probability parameter to an initial value; and

determining, for the first document at the value of the parameter, a probability of an assignment of the first document to the first topic cluster based on a word included in the first document and the first class.

18. The method of claim 15, wherein associating the first document with a second topic cluster comprises:

assigning the plurality of multi-word documents to a first class;

setting a probability parameter to an initial value; and

determining a probability of an assignment of the first document to the second topic cluster based on a word included in the first document and the first class.

19. The method of claim 15, wherein providing a representation comprises:
providing the representation after each document in the plurality of multi-word documents has been associated with at least one topic cluster corresponding to a sub-node in the hierarchy, wherein any of the plurality of multi-word documents may be

associated to more than one topic cluster based on different portions of the respective document.

20. A computer-implemented method for clustering data reflecting users, represented as a set of (data, user) pairs, into a hierarchical data structure including a root node associated with a plurality of sub-nodes, wherein each sub-node represents an action that is performed on a document collection, comprising:

accessing a user data collection reflecting a plurality of users who each perform at least one action on the document collection, wherein each action may be unique;

performing a clustering process that creates the hierarchical data structure, wherein the clustering processing comprises:

retrieving a first user data, associated with a first user, from the user data collection,

associating the first user data with a first sub-node based on a first action performed by the first user on the document collection, and

associating the first user data with a second sub-node provided the first user data is based on a second action, wherein the first and second sub-nodes are associated with different descendent paths of the hierarchical data structure;

storing a representation of the hierarchical data structure in a memory; and

making the representation available to an entity in response to a request associated with the user data collection.

21. The method of claim 20, wherein each action in the one or more actions includes:

writing to, printing, and browsing the document collection.

22. A computer-implemented method for clustering a plurality of images based on text associated with the images, where each image is represented as a set of pairs (image, image feature) and (image, text feature), into a hierarchical data structure including a root node associated with a plurality of sub-nodes, wherein each sub-node represents a different topic, the method comprising:

accessing an image collection;

performing a clustering process that creates the hierarchical data structure, wherein the clustering processing comprises:

associating a first image with a first sub-node based on a first portion of text associated with the first image, and

associating the first image with a second sub-node based on a second portion of text associated with the first image, wherein the first and second sub-nodes are associated with different descendant paths of the hierarchical data structure;

storing a representation of the hierarchical data structure in a memory; and

making the representation available to an entity in response to a request associated with the image collection.

23. A computer-implemented method for clustering customer purchases, represented as a set of (customer, purchase) pairs, into a hierarchical data structure including a root node associated with a plurality of sub-nodes, wherein each sub-node represents a group of customers who purchased the same type of product from one or more business entities, the method comprising:

accessing information associated with a plurality of customers who purchased various types of products from a plurality of business entities;

performing a clustering process that creates the hierarchical data structure, wherein the clustering processing comprises:

associating a first customer with a first sub-node based on a first type of product purchased from a first business entity, and

associating the first customer with a second sub-node provided the first customer is based on a second type of product that the first customer purchased from a second business entity, wherein the first and second sub-nodes are associated with different descendant paths of the hierarchical data structure;

storing a representation of the hierarchical data structure in a memory; and

making the representation available in response to a request associated with the customer data collection.

24. The method of claim 1, wherein the representation defines the probability of a document as the product of the probability of the (document, word) pairs it contains.

25. The method of claim 24, wherein the product is calculated after mixing the document-word pairs over the clusters.

26. The method of claim 25, wherein mixing the (document, word) pairs over the clusters comprises a probability model of the form:

$$P(x) = \sum_c P(c)P(x | c)$$

wherein c is the group of clusters involved in the calculation, and x is a (document, word) pair.

Evidence Appendix

None.

Related Proceedings Appendix

None.